

IN THE SPECIFICATION:

Please replace the paragraph beginning at page 13, line 7 in the specification with the following replacement paragraph:

— In accordance with the illustrative embodiment of the present invention, the cluster connection manager 375 is initialized early in the booting process of the storage operating system 300. Once initialized, the cluster connection manager initiates a peer-to-peer connection with its cluster partner as described further below. —

Please replace the paragraph beginning at page 13, line 20 in the specification with the following replacement paragraph:

— In accordance with an illustrative embodiment of the present invention, the cluster connection manager 375 of the storage operating system 300 performs all peer-to-peer communications between the storage system of cluster 130. To that end, the cluster connection manager 375 350 initially creates a peer-to-peer connection with the storage system's cluster partner using the novel method described below. It should be noted that while this description is written in terms of storage operating systems executing on clustered storage appliances, the principles of the present invention may be applied to any cluster of computers utilizing VI, InfiniBand or other protocols supporting RDMA operations. —

Please replace the paragraph beginning at page 14, line 26 in the specification with the following replacement paragraph:

— While in state 3 (Waiting on Partner), the cluster connection manager periodically examines a specified variable stored at a specific memory location. This variable is modified via a RDMA write operation by the cluster partner when the peer cluster connection manger reaches a either state 5 or the completed state. If the variable

is modified to alert the cluster connection manager that the partner is in state 5, then the cluster connection manager transitions to state 5 (event 435). Similarly, if the variable is modified so that it identifies the partner system as being complete with the initialization routine, the connection manager transitions to the connected states (event 440). —

Please replace the paragraph beginning at page 15, line 9 in the specification with the following replacement paragraph:

— If the RDMA read operation (action 410) succeeds, the cluster connection manager transitions to state 2 (Attempting Write) where it performs an RDMA write operation (action 435 415) to the storage system cluster partner. This RDMA write operation is directed to a specified and pre-determined memory location in the storage system partner, which is illustratively the same memory location that is monitored by the partner in state 3 (Waiting on Partner). If the RDMA write operation fails, the cluster connection manager moves to the error state 425. The failure of the RDMA write operation signifies that the storage system cluster partner is not in a ready state or that an error condition has occurred between the completion of the successful RDMA read operation and the execution of the RDMA write operation. —

Please replace the paragraph beginning at page 15, line 19 in the specification with the following replacement paragraph:

— However, if the write operation succeeds, the connection manager then attempts a second RDMA read operation (event 460) and transitions to state 4 (Attempting Read). If the read operation fails, the connection manager moves to the error state (event 465). Otherwise, the connection manager attempts an additional RDMA write operation (action 490) and transitions to state 5 (Attempting Write). Once in state 5, the connection manager will wait until its partner has reached state 5 (event 470 455)

before being connected. The connection manager will be alerted that the partner is in state 5 by the successful RDMA write operation from its partner. Once the connection manager is in the connected state, the cluster connection managers have established a connection with their corresponding peer connection manager. —

Please replace the paragraph beginning at page 15, line 29 in the specification with the following replacement paragraph:

— To again summarize, the present invention is directed to a novel system and method for establishing reliable peer communication in a clustered environment. Specifically, the technique enables a cluster pair to establish VI/QP connections without requiring upper level layers of the storage operating system to be fully initialized. The initial peer connection is established by performing a series of RDMA read and write operations. Once a reliable VI connection has been made, the cluster connection manager gathers connection information from cluster connection clients to be passed to its partner. The cluster connection manager then generates the appropriate VIs and alerts its partner that it has reached a ready state. Cluster connection clients may then utilize the created Vis for communicating with their cluster partner peer pair. —